

TITLE OF THE INVENTION

LAYER 2/LAYER 3 INTERWORKING VIA PHYSICAL LOOPBACK

INVENTOR
Edward SIERECKI

P24418.S03

LAYER 2/LAYER 3 INTERWORKING VIA PHYSICAL LOOPBACK

CROSS REFERENCE TO RELATED APPLICATION

[0001] The present application is related to a U.S. patent application having attorney docket number P24392 in the names of K. LIU et al., filed on November 12, 2003, the disclosure of which is expressly incorporated by reference herein in its entirety.

BACKGROUND OF THE INVENTION

Field of the Invention

[0002] The present invention relates to the field of communications. More particularly, the present invention relates to improving reliability when adding layer three routing to layer two networks.

Background Information

[0003] Figure 1 illustrates an example of today's networking environment in which layer two networks are provided with layer three routing. In today's networking environment, a customer edge device 10 connects to a layer two switch 12, such as an ATM switch or a frame relay switch. The switches 12 are interconnected with interoffice trunks 14. The connection 16 between the switch and the customer edge device 10 can be any known interface.

[0004] In an asynchronous transfer mode (ATM) example, a permanent virtual connection (PVC) 18 is configured from the ATM switch port connecting to the customer edge device 10 to a trunk 15 terminating at the far end of the switch 12. The trunk 15 is similar to the other trunks 14 shown, except, the trunk 15 is partitioned. The partitioning is required to separate a user-network interface (UNI) e.g., an Internet protocol (IP) interface, from the standard layer two trunk group. At least one partition is required for the standard ATM trunks, and another partition is required for each IP interface that is

defined. Complex provisioning and associated administrative burden are required to partition the trunk 15.

[0005] The IP interface across the trunk 15 is defined between the switch 12 and a platform 20, such as an Alcatel 7670 RSP (routing switch platform), available from Compagnie Financiere Alcatel of France. The defined IP interface on the ATM trunk 15 uses standard ATM encapsulation. The defined IP interface must also specify a virtual path identifier/virtual channel identifier (VPI/VCI) of the PVC connection 18, associated with the IP service, to the customer edge device 10.

[0006] The platform 20 includes layer two switching capabilities and layer three switching capabilities. In today's multi service platforms 20, the layer two portion is independent from and isolated from the layer three portion. Typically, a layer three port 22 of the multi service platform 20 terminates the UNI connection 15.

[0007] A problem associated with the current configuration is that when the UNI connection 15 (either the link or a port) fails, the layer two network will not re-route a circuit to the multi service platform 20 because the layer two network only extends to the connection 15. In other words, the PVC 18 terminates on the layer two switch 12, and not on the platform 20. Thus, no layer two protection is available for the trunk 15 between the switch 12 and the platform 20. If the trunk 15 fails, all customers using the link 15 would be out of service.

[0008] Current solutions addressing the single point of failure problem include dual homing from a customer site 10 to two different platforms 20. In this case, when one connection fails, the other connection can maintain connectivity. This approach, however, consumes too many network resources by requiring both paths to be permanently maintained, adding significant complexity to the provisioning and maintenance procedures for this service.

[0009] Another solution reduces the length of the connection between the switch 12 and the platform 20 by deploying the switches 12 and platforms 20 within the same central office. Thus, the connection 15 becomes an intra-central office connection. This solution, however, increases the overall switch deployment cost and is still subject to a single point of failure.

[0010] Thus, a solution is needed to address the single point of failure problem and the complexity of provisioning problem without increasing consumption of network resources.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The present invention is further described in the detailed description that follows, by reference to the noted drawings by way of non-limiting examples of embodiments of the present invention, in which like reference numerals represent similar parts throughout several views of the drawings, and in which:

Fig. 1 is a diagram showing a prior art networking environment; and

Fig. 2 is a diagram showing a networking environment, according to an aspect of the present invention.

DETAILED DESCRIPTION OF EMBODIMENTS

[0012] The present invention relates to increasing reliability of interconnected layer two and layer three networks. The increased reliability is achieved by providing a physical loopback between layer two and layer three switching components within a multi service platform.

[0013] In view of the above, the present invention through one or more of its various aspects and/or embodiments is presented to accomplish one or more objectives and advantages, such as those noted below.

[0014] According to an aspect of the present invention, a multi service platform includes a layer two switching component, a layer three switching component, and a physical loopback. The physical loopback connects the layer two switching component and the layer three switching component. The layer two capabilities and layer three capabilities are, therefore, integrated together. The physical loopback may be a fiber jumper cable.

[0015] In one embodiment, the layer two switching component and the layer three switching component are on a line card. Both ends of the loopback may terminate on the line card. At least one additional physical loopback may be provided, connecting to another layer three switching component on the line card. Thus, redundancy for the layer three functionality is provided on the line card. At least one additional line card may include another layer two switching component and another layer three switching component. Thus, the at least one additional line card provides redundancy. The additional line card(s) may include at least one additional physical loopback terminating on the additional line card(s).

[0016] According to another aspect of the present invention, a network includes multiple layer two switches, and at least one platform. The platform includes a layer two switching component, a layer three switching component and a physical loopback between the layer two switching component and the layer three switching component. The network also includes at least one connection between one of the layer two switches, which communicates with a customer edge device, and the layer two switching

component of the platform. Thus, a failure of the connection, which extends to the platform, is protected by layer two network failure restoration.

[0017] The layer two network switches may be ATM switches. Moreover, the connection(s) may be a permanent virtual connection (PVC). Further, the layer two switching component of the platform may be an ATM switch, and the layer three switching component of the platform may be an IP router.

[0018] In yet another aspect, a method is provided for routing traffic across a layer two network having layer three routing capabilities. The method includes routing traffic from a customer across the layer two network to a layer two switching component in a platform, and routing traffic from the layer two switching component across a physical loopback to a layer three switching component in the platform. The method also includes determining, at the layer three switching component, where to route the traffic, returning the traffic to the layer two switching component, and forwarding the traffic to a destination based upon the determined route. In one embodiment, the layer two network is an ATM network.

[0019] The various aspects and embodiments of the present invention are described in detail below.

[0020] The present invention improves reliability of layer two networks having layer three routing by extending the layer two network to a layer two switching component of a multi service platform. Thus, if an interface between the layer two switch and the layer two switching component of the multi service platform fails, the layer two network failure recovery scheme re-routes the circuit to the layer two portion of the multi service platform.

[0021] Referring now to figure 2, an embodiment of the present invention is shown. The multi service platform 20 includes a physical loopback 30 which connects the layer

two switching component 24 of the platform 20 and a layer three switching component 22 of the platform 20. In one embodiment, the loopback 30 is an OC3 or OC12 fiber jumper cable. The cable may be approximately two feet long. Of course the loopback 30 is not limited to this length and is also not limited to the OC3 or OC12 throughput values. The layer two and layer three portions 22, 24 may be provided on a line card, e.g., an Alcatel MR8, within the platform 20.

[0022] In one embodiment, the layer two network is an ATM network, and the layer three network is an IP network, although any other type of layer two and layer three networks can be provided, for example, ethernet, frame relay and multiprotocol label switching (MPLS). In the ATM/IP embodiment, the multi service platform 20 includes an ATM switch as the layer two portion 24 and an IP router as the layer three portion 22.

[0023] In the ATM/IP embodiment, the ATM network terminates on the ATM switching component 24 in the platform 20. That is, a PVC 18 connects all the way to the ATM port 24. Although a PVC is described in this example, any other type of ATM connection, e.g., a switched virtual connection (SVC) can be used. The physical loopback 30 connects the IP component 22 and the ATM component 24. Thus, the ATM network includes the ATM switch 24 on the platform 20 and accordingly protects against ATM trunk 14 failures for all trunks terminating on the switch with its failure restoration.

[0024] The physical loopback 30 thus becomes the only unprotected link. The failure probability of the link 30 is low, however, because the link is so short. Moreover, the reliability is further enhanced by the extension of the layer two network recovery scheme.

[0025] When data arrives at the ATM port 24 from the PVC 18, the cells are forwarded to the IP interface 22. The IP interface then performs an IP look-up to determine the destination of the traffic. The traffic is then returned to the ATM component 24 and ultimately to its destination 32 via, e.g., another PVC 34.

[0026] In one embodiment, both ends of the loopback 30 are on the same line card. Thus, the IP functionality is isolated to a single card, permitting protection with line card redundancy. In a further embodiment, redundant physical loopbacks are provided. That is, each card can be provided with a physical loopback 30. Line card redundancy could also be implemented on a single line card for each platform to provide redundancy for the layer three functionality. Redundancy on line cards terminating layer two trunks is less critical because layer two re-routing occurs if these cards fail.

[0027] An advantage of the present invention is that all UNIs can be provisioned to the same layer two switching component 24. Then, the loopback 30 can extend from the port 24 to the IP interface 22. Accordingly, end-to-end ATM connections can be provided. In other words, the ATM trunks 14 remain as pure ATM connections, obviating the need to partition the trunks 14 to provide separate IP interfaces. Consequently, existing CACing and bandwidth engineering methods can still be used. Moreover, simplified provisioning and better ATM bandwidth utilization occur.

[0028] Thus, the present invention provides a physical loopback connection between layer two and layer three switching components of a multi service platform thereby improving end to end reliability. It is understood that the words that have been used are words of description and illustration, rather than words of limitation. Changes may be made within the purview of the appended claims, as presently stated and as amended, without departing from the scope and spirit of the invention in its aspects. Although the invention has been described with reference to particular means, materials and embodiments, the invention is not intended to be limited to the particulars disclosed; rather, the invention extends to all functionally equivalent structures, methods, and uses such as are within the scope of the appended claims.

[0029] In accordance with various embodiments of the present invention, the methods described herein are intended for operation as software programs running on a computer processor. Dedicated hardware implementations including, but not limited to, application specific integrated circuits, programmable logic arrays and other hardware devices can likewise be constructed to implement the methods described herein. Furthermore, alternative software implementations including, but not limited to, distributed processing or component/object distributed processing, parallel processing, or virtual machine processing can also be constructed to implement the methods described herein.

[0030] It should also be noted that the software implementations of the present invention as described herein are optionally stored on a tangible storage medium, such as: a magnetic medium such as a disk or tape; a magneto-optical or optical medium such as a disk; or a solid state medium such as a memory card or other package that houses one or more read-only (non-volatile) memories, random access memories, or other re-writable (volatile) memories. Accordingly, the invention is considered to include a tangible storage medium or distribution medium, as listed herein and including art-recognized equivalents and successor media, in which the software implementations herein are stored.

[0031] Although the present specification describes components and functions implemented in the embodiments with reference to particular standards and protocols, the invention is not limited to such standards and protocols. Each of the standards for layer two and layer three transmission represent examples of the state of the art. Such standards are periodically superseded by faster or more efficient equivalents having essentially the same functions. Accordingly, replacement standards and protocols having the same functions are considered equivalents.